

Dark halo mergers and the formation of a universal profile

D. Syer and Simon D.M. White

Max-Planck-Institut für Astrophysik, 85740 Garching-bei-München, Germany.

ABSTRACT

We argue that a universal density profile for dark matter halos arises as a natural consequence of hierarchical structure formation. It is a fixed point in the process of repeated mergers. We present analytic and numerical arguments for the emergence of a particular form of the profile. At small radii, the density should vary as $r^{-\alpha}$, with α determined by the way the characteristic density of halos scales with their mass. If small halos are dense, then α is large. The mass-density relation can be related to the power spectrum of initial fluctuations, $P(k)$, through ‘formation time’ arguments. Early structure formation leads to steep cusps. For $P(k) \sim k^n$ we find $\alpha \simeq 3(3+n)/(5+n)$. The universal profile is generated by tidal stripping of small halos as they merge with larger objects.

Key words: cosmology: theory – dark matter – stellar dynamics – galaxies: kinematics and dynamics

1 INTRODUCTION

In a series of recent papers, Navarro, Frenk & White (1995, 1996, 1997, hereafter NFW) have claimed that dark halos formed by dissipationless hierarchical clustering from gaussian initial conditions can be fitted by a universal density profile of the form

$$\rho(r) = \delta \frac{R^3}{r(R+r)^2}, \quad (1)$$

where R is a characteristic length scale and δ a characteristic density. This function fits the numerical data presented by NFW over a radius range of about two orders of magnitude. Equally good fits are obtained for high mass (rich galaxy cluster) and for low mass (dwarf galaxy) halos, and in cosmological models with a wide range of initial power spectra $P(k)$, density parameters Ω , and cosmological constants Λ . In any given cosmology the simulated halos show a strong correlation between R and δ ; low mass halos are denser than high mass halos. NFW interpret this as reflecting the fact that low mass halos typically form earlier. In independent work Cole & Lacey (1996) came to similar conclusions based on studies of a series of cosmological simulations with $P \propto k^n$ and $\Omega = 1$. We shall refer to the profile of equation (1) as the NFW profile.

It will prove useful below to consider a broader family of density profiles: we will adopt the family defined by

$$\rho(r) = \delta \frac{R^{\alpha+\beta}}{r^\alpha (R^{1/\gamma} + r^{1/\gamma})^{\gamma(\beta-\alpha)}}. \quad (2)$$

This family has been studied extensively by Zhao (1996); subfamilies were studied by Dehnen (1995) and by Tremaine *et al.* (1994). The NFW profile is of this form: an inner cusp

with logarithmic slope $\alpha = 1$, an outer envelope with logarithmic slope $\beta = 3$, and a ‘turn over exponent’ $\gamma = 1$. Another well-known example is the Hernquist (1990) profile: $(\alpha, \beta, \gamma) = (1, 4, 1)$. Both NFW and Cole & Lacey (1996) compared the Hernquist profile to their numerical data. Although it fits high mass halos almost as well as the NFW profile, its more rapid fall-off does not fit low mass halos adequately. Detailed comparisons of both profiles with high resolution simulations of galaxy clusters are also described by Tormen, Bouchet & White (1996).

In this paper we argue that a universal profile arises as a result of repeated mergers. Violent relaxation of a finite mass *isolated* system leads to a profile with $\rho \propto r^{-4}$ as r tends to infinity (Aguilar & White 1986, Jaffe 1987, Merritt, Tremaine & Johnstone 1989, Barnes & Hernquist 1991). Sufficiently strong violent relaxation, either through repeated merging or through cold inhomogeneous collapse, produces $\rho \propto r^{-3}$ over the regions which contain most of the mass (White 1979, Villumsen 1982, Duncan, Farouki & Shapiro 1983, McGlynn 1984). There are, however, good reasons why neither of these results should apply to halo formation through hierarchical clustering. Cosmological mergers are not isolated, they occur continually and from bound orbits, and they usually involve objects of very different mass. These properties are reflected in the other simple paradigm for halo formation, the spherical infall model (Gunn & Gott 1972, Fillmore & Goldreich 1984, Hoffman & Shaham 1985, White & Zaritsky 1992). This model predicts relatively shallow density profiles, $\rho \sim r^{-2}$ for the virialised regions of halos.

The next section uses simple arguments to show why merging in a cosmological context might give rise to a characteristic density profile with a central cusp. The third sec-

tion then carries out some numerical experiments to test these arguments. A final section discusses our results.

2 MERGERS AND TIDAL DISRUPTION

Hierarchical clustering proceeds through a continual series of unequal mergers. Small objects typically form at earlier times and so have higher characteristic densities. The results of Navarro, Frenk & White (1996) for a CDM universe can be fit by a power law

$$\delta \propto M^{-\nu} \quad (3)$$

with $\nu \simeq 0.33$. The characteristic radius is given by

$$R \propto M^{(1+\nu)/3}. \quad (4)$$

Simple scaling arguments using linear theory (Kaiser 1986) predict that if $P(k) \sim k^n$ then

$$\nu = (3+n)/2. \quad (5)$$

$\nu = 0.33$ corresponds to $n = -2.3$ as expected for a CDM universe on the relevant scales.

Consider a large parent halo merging with a smaller satellite system. Two important dynamical processes are responsible for the evolution of the system. Dynamical friction causes the satellite's orbit to decay, bringing it down into the central regions. Tidal stripping removes material from the satellite and adds it to the diffuse mass of the parent. Our claim is that the combination of these processes leads to $\alpha \simeq 1$ in the inner cusp. This structure is a stable fixed point in the process of repeated mergers, in a sense to be defined below.

A crude but useful way of understanding tidal stripping, going back at least to the work of van Hörner (1957), is as follows. Consider a satellite of mass m orbiting a parent of mass M at mean distance r . The frequency Ω of the orbit is given roughly by $\Omega^2 \sim \bar{\rho}_M(r)$, where $\bar{\rho}_M(r)$ is the average density of the parent interior to r . Next consider a particle Q orbiting within the satellite at a mean distance ξ from its centre. The frequency ω of this orbit is given by $\omega^2 \sim \bar{\rho}_m(\xi)$, where $\bar{\rho}_m(\xi)$ is the average density of the satellite interior to ξ . If $\Omega = \omega$ there is a resonance between the force the satellite exerts on Q, and the tidal force exerted by the parent. This results in a secular transfer of energy to Q, which then escapes from the satellite. Thus, crudely speaking, the satellite is stripped of material down to radius ξ defined by $\omega(\xi) = \Omega(r)$. The newly stripped material will orbit throughout the parent but will, on average, remain near the orbit of the satellite at the instant that it was stripped.

Weinberg (1994abc, 1996) has shown that material can actually be stripped off the satellite substantially within the primary resonance $\Omega = \omega$. The strongest effect is nevertheless at this resonance, and since the timescale of real cosmological mergers is not very different from Ω^{-1} , there is little time for weaker effects to be felt. For simplicity we stick to the standard assumption of stripping at the primary resonance, and we will use N -body experiments to check our principal results.

Since clustering is hierarchical we expect the satellite to have a higher characteristic density and a smaller characteristic radius than the parent. As a result the density of the

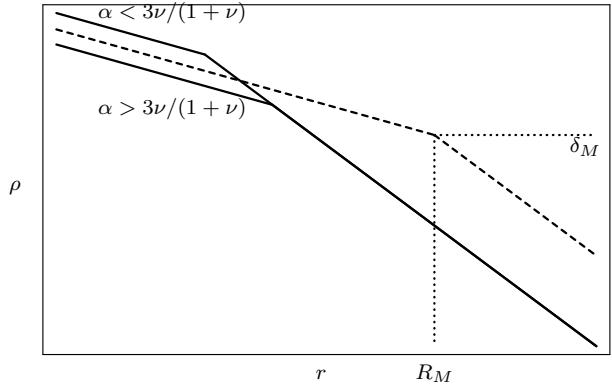


Figure 1. The superimposed density profiles of a parent halo of mass M and of homologous satellite halos of mass $m < M$. The logarithmic slope of the central cusp is $-\alpha$. Two satellites are shown, one with $\rho_m(R_m) < \delta_m$ (lower solid curve) and one with $\rho_m(R_m) > \delta_m$ (upper solid curve). The former has $3\nu/(1+\nu) < \alpha$ and the latter $3\nu/(1+\nu) > \alpha$.

parent at the characteristic radius of the satellite $\rho_M(R_m)$ can be either larger or smaller than the density of the satellite at the same radius $\rho_m(R_m) = \delta_m$ (see Figure 1). (For $\alpha < 3$ the average density is also proportional to $r^{-\alpha}$.) Thus, when dynamical friction has brought the satellite right to the centre of the parent, it will either have been entirely disrupted, or will have survived largely intact; the result depends on whether $\rho_M(R_m)$ is less than or greater than δ_m . The case of equality, where the satellite is marginally disrupted by the parent as it reaches the centre, is given by

$$\rho_M(R_m) = \delta_m \left(\frac{R_M}{R_m} \right)^\alpha = \delta_m, \quad (6)$$

which can be written in terms of the masses using equations (4) and (3) as follows

$$1 = \left(\frac{M}{m} \right)^{\nu - \alpha(1+\nu)/3}, \quad (7)$$

whence

$$\alpha = \frac{3\nu}{1 + \nu}. \quad (8)$$

In conjunction with (5) this predicts

$$\alpha = 3 \left(\frac{3+n}{5+n} \right). \quad (9)$$

It is interesting that this exponent is the same as that predicted in a completely different context, but for similar reasons, for the small scale behaviour of mass autocorrelation function in hierarchical clustering on the stable clustering hypothesis (Davis & Peebles 1977). These equations constitute our prediction for the logarithmic slope of the inner cusp in cosmological halos. For the Navarro, Frenk & White (1996) halos, $\nu \simeq 0.33$ implying $\alpha \simeq 0.75$. This is probably consistent with NFW's fits to these halos since their experiments do not resolve the central cusp to $r \ll R$. In fact α does not vary strongly with n for cosmologically interesting values of -1 to -2.5 . The simulations of Cole & Lacey (1996) include the cases $n = \{0, -1, -2\}$ corresponding to $\alpha = \{1.8, 1.5, 1\}$. They do not resolve the central cusps well,

but there is a hint in their Figure 9 that $n = 0$ has steeper halos than $n = -2$.

A further argument is needed to show that halos will actually prefer the special value $\alpha = 3\nu/(1 + \nu)$. Suppose α in a parent halo were smaller than this value, and it merged with a homologous satellite. In this case $\rho_M(R_m) < \delta_m$, so the satellite survives intact right to the centre. It thus boosts the density of the parent cusp in the region $r < R_m$, and steepens the effective logarithmic slope in $r < R_M$; i.e. α increases. Now suppose α in the parent were larger than $3\nu/(1 + \nu)$. In this case $\rho_M(R_m) > \delta_m$, so the satellite will be disrupted at $r > R_m$, spreading its mass out over a region which is larger than its own characteristic radius. It thus boosts the density of the cusp in the region $r > R_m$, and softens the effective logarithmic slope in $r < R_M$. Thus hierarchical halo formation by repeated mergers should lead to the density profile taking on a fixed point form, which depends on the mass-density relation of merging subunits.

3 NUMERICAL EXPERIMENTS

We now test the assertion that there is a fixed point in the process of repeated mergers. In outline our procedure is: to merge a halo repeatedly with a small satellite that is created by scaling the parent according to equation (3). We emphasise that the satellite is always a copy of the *current* halo, i.e. the latest merger product.

First we take a parent halo with mass M but with an initial density profile quite different from the NFW profile:

$$\rho_0(r) = \delta \frac{R^\beta}{(R^2 + r^2)^{\beta/2}}, \quad (10)$$

with $\delta = 1$, $R = 1$ and β an input parameter. The following steps are then carried out repeatedly:

- 1 Choose a mass ratio for the next merger, $f = m/M$, by taking a random value from a distribution between $f_{\min} = 10^{-3/(1+\nu)}$ and $f_{\max} = 0.4$ in which the number of daughters with mass ratios in $(f, f + df)$ is specified by

$$dN = f^\eta df. \quad (11)$$

Thus $\eta = 0$ gives a uniform distribution of daughter masses, and $\eta < 0$ produces a larger number of low mass daughters.

- 2 Create a satellite of mass fM by cloning the parent homologously: $\rho \rightarrow \rho f^{-\nu}$, $r \rightarrow r f^{(1+\nu)/3}$.
- 3 Merge the satellite with the parent. Section 3.1 describes a semi-analytic model of the merging process based on the arguments in Section 2. Section 3.2 describes an N -body implementation of the merging process.

3.1 Semi-analytic models

First we implement our simplified picture of tidal stripping as described in Section 2. On a logarithmic grid in r between $r = 0.01$ and $r = 100$, a merger is simulated as follows. Work inwards from the outer edge of the parent, removing all the mass from the satellite where $\rho_m(\xi) = \rho_M(r)$ and placing it in the radial bin at r . If the profiles cross ($\rho_M(R_m) < \delta_m$) add the satellite bin by bin to the parent, as if it survived intact to the centre.

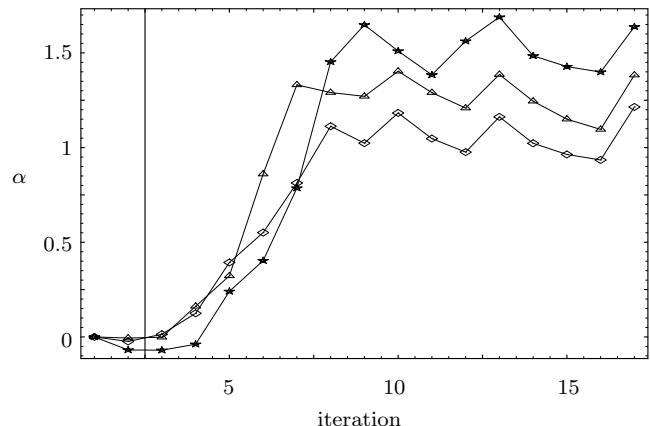


Figure 2. The logarithmic slope of the cusp α as a function of the number of mergers in the experiments of Section 3.1. From bottom to top at the right $\nu = (0.33, 0.5, 0.75)$.

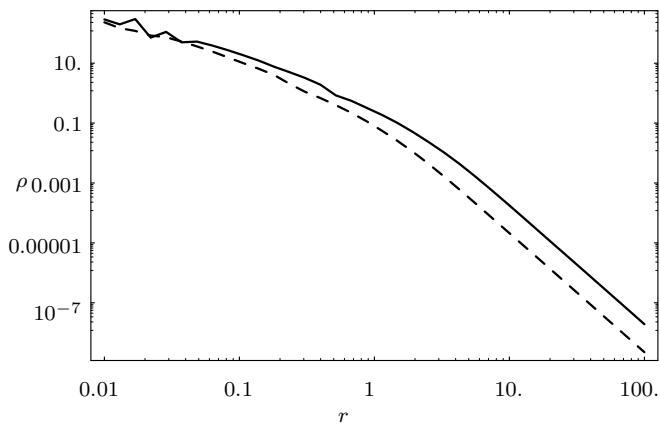


Figure 3. The parent halo profile from the last iteration in the experiment with $(\nu, \eta, \beta) = (0.33, 0, 4)$ superimposed on a homologous satellite with mass ratio $f = 0.22$. Compare Figure 1.

Figure 2 shows the evolution of the profile as the iterations proceed for runs with $(\eta, \beta) = (0, 4)$ and $\nu = (0.33, 0.5, 0.75)$. Clearly the homogeneous core is steepened and settles down into a cusp, as it was argued it should in the previous section. A steady state is reached, which, as predicted, automatically flattens itself if it gets too steep. The final value of α , averaged over 10 iterations, is summarised in Table 1 and in Figure 4. As shown by Figure 4, the value of α is systematically larger than predicted by equation (8). Figure 4 shows the results of experiments with $\beta = 3$ and $\beta = 4$. The lower value of β tends to produce slightly lower values of α , but still consistent with the $\beta = 4$ results at the 1σ level.

Figure 3 shows the profile of a satellite superimposed on its parent, so it is the equivalent of Figure 1. Since α is slightly greater than predicted by equation (8), the satellite is less dense than its parent at small radii. Note that the inner cusp is very close to a power law. However, Figure 5 shows that the slope α is correlated with the turn over exponent γ . This could mean that we are not resolving the power law in the cusp. To test this we compare the measured value of α with the *actual* logarithmic slope of the profile at the

Table 1. The results of the experiments of Section 3.1. The input parameters are (β, η, ν) , and the output fitted profiles are specified by α and γ . The errors in the last two decimal places are given in brackets (1σ estimated from the last 10 iterations in each case). $\alpha_{\text{eff}}/\alpha$ is a measure of how well we resolve the inner cusp (see text).

β	η	ν	α	γ	$\alpha_{\text{eff}}/\alpha - 1$ (%)
3	0	0.05	0.46(10)	0.57(08)	0.4
3	0	0.125	0.62(08)	0.64(12)	0.7
3	0	0.2	0.76(08)	0.70(13)	0.9
3	0	0.33	1.00(13)	0.75(12)	0.9
3	0	0.5	1.22(14)	0.80(16)	1.0
3	0	0.75	1.47(14)	0.79(20)	0.9
4	0	0.05	0.59(09)	0.54(03)	0.1
4	0	0.125	0.81(10)	0.65(05)	0.3
4	0	0.33	1.06(10)	0.83(13)	1.4
4	0	0.5	1.27(10)	0.88(16)	1.5
4	0	0.75	1.52(11)	0.98(16)	1.7
4	-0.8	0.33	1.08(07)	0.73(07)	0.6
4	-0.7	0.33	1.14(08)	0.75(05)	0.6
4	-0.5	0.33	1.10(09)	0.87(06)	1.3
4	0.5	0.33	1.00(09)	0.95(15)	2.9
4	1.0	0.33	1.10(06)	0.76(09)	0.7
4	2.0	0.33	1.12(08)	0.68(09)	0.4

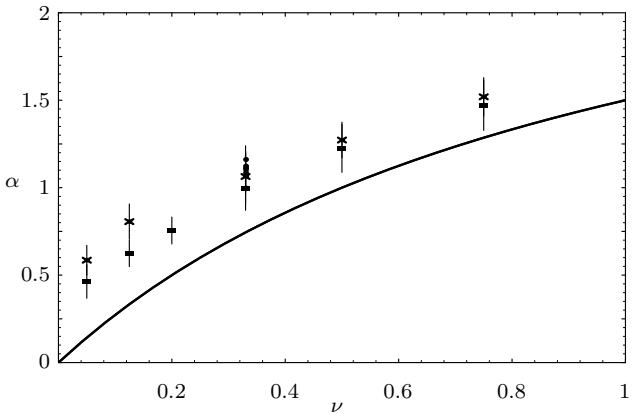


Figure 4. The relationship between ν and α as measured (with 1σ error bars). Times signs are for $\beta = 4$, squares are for $\beta = 3$ (both with $\eta = 0$) and no symbol for $\beta = 4$ and different values of η (see Table 1). The solid line is the prediction of equation (8).

inner grid point:

$$\alpha_{\text{eff}} = -\frac{d \ln \rho}{d \ln r} = \frac{\alpha + \beta r^{1/\gamma}}{1 + r^{1/\gamma}}. \quad (12)$$

The last column of Table 1 shows that the difference between α and α_{eff} is small, thus we have adequately resolved the central power law.

Calculations were performed with $(\beta, \nu) = (4, 0.33)$ and f drawn from different distributions. We adjust the relative numbers of large and small mergers by changing the value of η in equation (11). The results are summarised in Table 1, and also appear in Figures 5 and 4. Generally the resulting values of α are similar, show no trend with η , and are statistically consistent with each other. There is nothing in the argument leading to equation (8) which depends on the mass ratio, so this is expected. The measured γ also shows no distinct trend with η , although, as can be seen from Figure 5, varying η can lead to large scatter in the value of γ .

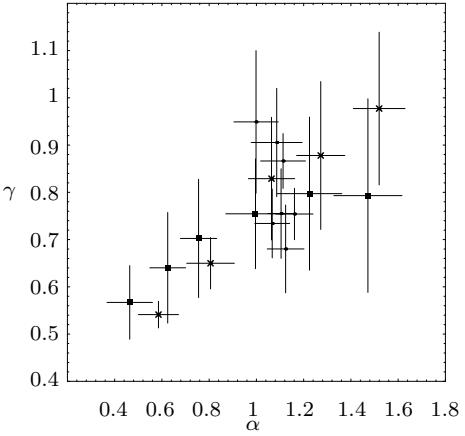


Figure 5. The correlation between γ and α (with 1σ error bars). Symbols as in Figure 4.

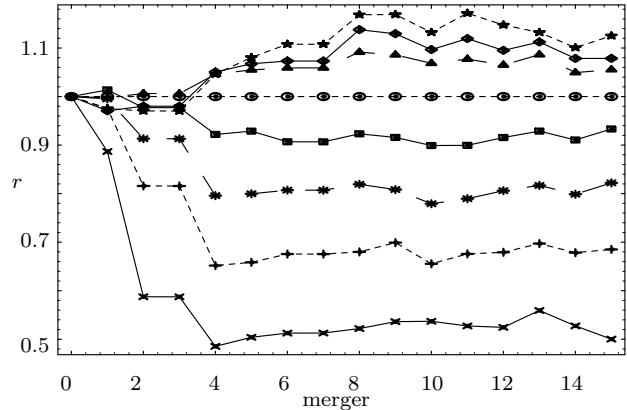


Figure 6. From bottom to top the ratio of the 10–50 and 80, 70, 60%-mass radii to the half-mass radius of the halo in simulation A as a function of the number of mergers. Each quantity is normalised to its initial value.

All these values are, however, consistent with each other at the 1σ level.

3.2 N-body experiments

The code we used was the publicly available adaptive P³M code of Couchman, Pearce and Thomas (1996) (configured to simulate isolated, collisionless systems). Compared with these N -body experiments, the simulations of Section 3.1 have better resolution and they are much cheaper to perform.

An initial model of a Plummer sphere ($\beta = 5$ in equation 10) is set up with 8000 particles. The satellite is placed at the 90% mass radius of the parent. Its orbit is such that it is at apocentre, and it has an angular momentum κ times that of the circular orbit with the same energy. For $\kappa = 0.3$, the apocentre (pericentre) of the satellite orbit is thus approximately $4R$ ($0.1R$), where R is the scale radius of the parent Plummer sphere. The satellite mass is drawn randomly according to equation (11) with $(f_{\min}, f_{\max}, \eta) = (.05, .2, 0)$. The different values of ν and κ used are summarised in Table 2. The satellite is produced by randomly eliminating a fraction $(1 - f)$ of the particles

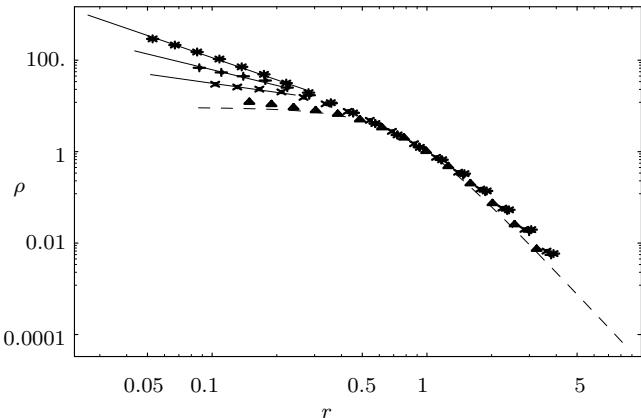


Figure 7. The average profile of the last 7 merger products of simulation A (crosses), B (plus signs) and C1 (stars) with a fit to the inner parts overlaid. The dashed line is a fit to the original profile measured in the same way (the fit parameters match a Plummer model well). The triangles close to the dotted line are a sanity check on the N -body code (see text).

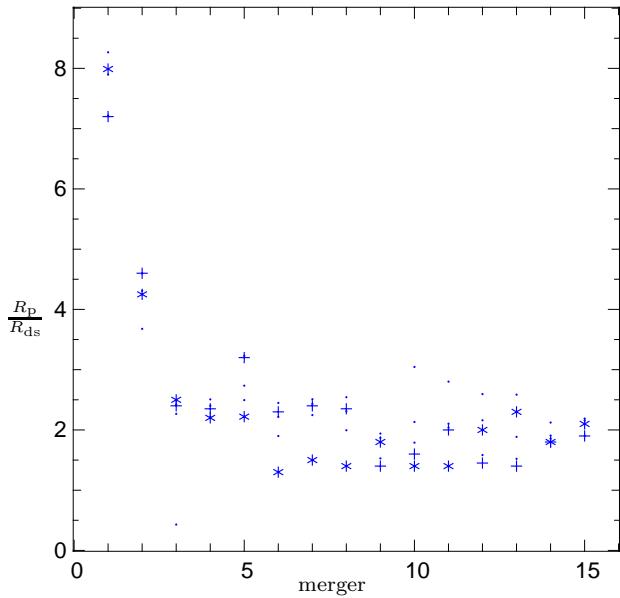


Figure 8. The ratio of $r_{1/2}$ in the parent halo to that of the disrupted satellite particles *after* the merger, as a function of the number of mergers, in simulations A (pluses), B (stars) and C[123] (dots).

of the parent and then scaling the positions and velocities of the remaining particles. The smoothing length used in the N -body code is a fixed fraction of the simulation box, which is not a fixed fraction of the halo scale length. The largest value it took in any of the simulations was one third of the radius of the innermost point at which the density was measured. In simulation A it was at most an eighth of this radius. Simulation A was also re-run with the smoothing length 10 times larger—the result was not significantly different.

After each merger another satellite is generated in the same way, and then merged. We measure the profile of the halo by starting at the potential centre and binning outwards

Table 2. The parameters of the N -body simulations in Section 3.2. The input parameters are (ν, κ) , and the output fitted profiles are specified by α . The last column is the value of α predicted by equation (8).

Model	ν	κ	α	$3\nu/(1+\nu)$
A	0.33	0.3	0.66	0.75
B	0.75	0.3	1.14	1.2
C1	1.5	0.3	1.75	1.8
C2	1.5	0.6	1.63	1.8
C3	1.5	0.05	1.79	1.8

in radius. The 10%, 20%,... radii are determined at the same time. Figure 6 shows the evolution of the ratios of these decile radii to the half-mass radius $r_{1/2}$. This shows that the structure converges after 7 or 8 mergers. Figure 7 shows the average profile of the last 7 mergers for models A, B and C1. Also shown is the initial Plummer sphere, and a sanity check: the profile of the initial model, evolved in isolation for the same time as all the mergers added together in simulation A. When evolved in isolation the profile changes a little due to numerical effects, but much less than when mergers are included. The cusp slope α was measured in these models by fitting a power law to the profile at $r < 0.3r_{1/2}$. The results are given in Table 2 and appear as solid lines in Figure 7. The general prediction that larger ν produces larger α is supported, and the actual values of the measured slopes are quite close to the prediction of equation (8).

A comparison of the three models with $\nu = 1.5$ (C1, C2, C3) shows that the profiles do not differ significantly. Thus the final profile is not sensitive to the value of κ .

Cosmological halos are not isolated, as our simulated halos are, and they are not in equilibrium in their outer parts. Thus we cannot make direct inferences about the outer profiles of cosmological halos on the basis of our simulations. Isolated merger products are expected to have an outer profile with $\beta = 4$ (Jaffe 1987, Merritt, Tremaine & Johnstone 1989). Thus in Figure 7 we see a slow roll-over from the inner cusp to an outer envelope with $\beta \simeq 4$, significantly less steep than the initial Plummer model. In cosmological halos, $\beta \simeq 3$ is observed near the virial radius for a wide range of models. It is not easy to identify the equivalent of the virial radius in our simulations, but the effective β of our profiles is about 3 in the range $(0.5, 2)r_{1/2}$. The 80% radius is at about $2r_{1/2}$, and the 90% radius (the apocentre of the satellite orbits) is at about $2.5r_{1/2}$. For comparison, in cosmological infall models the turnaround radius is at about three times the virial radius.

We also test the ‘sinking satellite’ assumption of Section 2, which was used in the calculations of Section 3.1. If dense satellites really sink to the centre of the parent then the particles which originated in the satellite should be concentrated towards the centre of the merger product. We test this by measuring the half-mass radius of the merger product and comparing it with that of the disrupted satellite particles. The result is shown in Figure 8. The degree of concentration of the disrupted satellite particles varies as the profile changes. The effect is most pronounced when the halos have a homogeneous core, and the satellite is much denser than the parent in its inner parts. As the cusp develops the satellite is disrupted at a progressively larger radius, as we argued should happen in Section 2, and the concentration settles down at a value of around 2 for all the simulations.

4 DISCUSSION

The analytic argument leading to equation (8) is highly idealised in its use of a single power law for the cusp profile. The semi-analytic model we have used for merging is based on the same idea but does not assume a form for the profile. Nevertheless it is probably not very accurate in detail. The *N*-body experiment gives us confidence that our simple arguments might be close to the truth. The values of α , measured from our experiments, agree quite well with the prediction of equation (8). Generally the semi-analytic experiments give α values higher than predicted, and the *N*-body experiments give lower values.

The physical origin we suggest for the cuspy profiles of halos is that sufficiently dense satellites can sink intact to the halo centre. The details of our merger prescription may affect our detailed conclusions, such as equation (8), but the robustness of the main physical idea should lead to similar behaviour in more realistic models. Figure 8 shows that the basic picture of Section 2 is not far wrong, at least for the kind of mergers studied here.

The extent to which this basic picture applies to actual cosmological halos remains to be demonstrated. Navarro, Frenk and White (1996) introduced the NFW profile for CDM halos. Navarro, Frenk and White (1997), in agreement with the lower resolution results of Cole and Lacey (1996), find that the NFW density profile continues to be a good fit to halos from power-law initial fluctuation spectra. The expected trend in the mass-density relation of halos is clearly visible in their fits: smaller effective ν for more negative n . According to the arguments of the present work, we would expect to see a corresponding trend in profile properties between $n = 0$ and $n = -1.5$. Nevertheless Navarro, Frenk and White (1997) are able to use a single fitting formula for all their halos.

The apparent contradiction may be accounted for in a number of ways. Cosmological halos always show a degree of scatter in their measured properties because of substructure—which is absent by design in our experiments. More data on the profiles of similar mass halos would give better statistical confidence in statements about trends in halo properties. In addition, it is difficult to find a consistent and dynamically meaningful scale for ρ and r on which to base a comparison of halos. In Figure 7 the trend between the profile and ν is obvious because all the halos line up so well in their outer parts. This in turn is a consequence of the fact that they are dynamically isolated, and hence the half-mass radius and density can be unambiguously defined. Since cosmological halos are not isolated, the same comparison cannot be made in this case. Next, the mass-density relation of halos in Navarro, Frenk and White (1997) has some scatter in it. If this is a consequence of a real scatter in the underlying relation it should produce a scatter in halo profile properties in addition to the effects of substructure. Finally, our experiments do not properly account for the expected systematic dependence on the power spectrum of the distributions of satellite mass and orbital parameters. Perhaps these could compensate in some way for the difference in core structure produced by the different mass-density relation. Since in our experiments we find no dependence of the profile on mass-distribution or orbital eccentricity, and since to exactly counter the effects of the mass-density rela-

tion would require fine tuning, we do not find this explanation compelling.

One thing is clear from our experiments though: it takes only a few mergers to establish the cusp profile. This means that the cusp can adjust to local conditions quite quickly and so should not depend in a complicated way on the *full* merger history of a halo.

REFERENCES

Aguilar, L., Ostriker, J.P., & Hut, P., 1988, ApJ, 335, 720
 Barnes, J.E., & Hernquist, L., 1991, ApJ, 370, L65
 Cole, S., & Lacey, C., 1996, MNRAS, in press, astro-ph/9510147
 Couchman, H.M.P., Pearce, F.R., & Thomas, P.A., 1996, astro-ph/9603116
 Davis, M., & Peebles, P.J.E., 1977, ApJ Supp., 34, 425
 Dehnen, W., 1995, MNRAS, 274, 919
 Dubinski, J., & Carlberg, R.G., 1991, ApJ, 378, 496
 Duncan, M.J., Farouki, R.T., & Shapiro, S.L., 1983, ApJ, 271, 22
 Fillmore, J.A., & Goldreich, P., 1984, ApJ, 281, 1
 Gunn, J.E., & Gott, J.R. III, 1972, ApJ, 176, 1
 Hernquist, L., 1990, ApJ, 356, 359
 Hoffman, Y., & Shaham, J., 1985, ApJ, 297, 16
 Jaffe, W., 1987, *Structure and Evolution of Elliptical Galaxies*, p.511, ed. de Zeeuw, T., Riedel, Dordrecht.
 Kaiser, N., 1986, MNRAS, 222, 323
 McGlynn, T., 1984, ApJ, 281, 13
 Merritt, D., Tremaine, S., & Johnstone, D., 1989, MNRAS, 236, 829
 Navarro, J.F., Frenk, C.S., & White, S.D.M., 1995, MNRAS, 275, 270
 Navarro, J.F., Frenk, C.S., & White, S.D.M., 1996, ApJ, 462, 563
 Navarro, J.F., Frenk, C.S., & White, S.D.M., 1997, in preparation.
 Tormen, G., Bouchet, F., & White, S.D.M., 1996, MNRAS, submitted, astro-ph 9603132
 Tremaine, S., et al, 1994, AJ, 107, 634
 van Hörner, S., 1957, ApJ, 125, 451
 Villumsen, J.V., 1982, MNRAS, 199, 493
 Weinberg, M.D., 1994a, AJ, 108, 1398
 Weinberg, M.D., 1994b, AJ, 108, 1403
 Weinberg, M.D., 1994c, AJ, 108, 1414
 Weinberg, M.D., 1996, astro-ph/9607099
 White, S.D.M., 1979, MNRAS, 189, 831
 White, S.D.M., & Zaritsky, D., 1992, ApJ, 394, 1
 Zhao, H.S., 1996, MNRAS, 278, 488

This paper has been produced using the Royal Astronomical Society/Blackwell Science TeX macros.